University of Groningen

Benchmarking UAQ solvers

Armando, Alessandro; Gazzarata, Giorgia; Turkmen, Fatih

Published in:
Proceedings of ACM Symposium on Access Control Models and Technologies, SACMAT

DOI:
[10.1145/3381991.3395616](https://doi.org/10.1145/3381991.3395616)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2020

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Armando, A., Gazzarata, G., & Turkmen, F. (2020). Benchmarking UAQ solvers. In *Proceedings of ACM Symposium on Access Control Models and Technologies, SACMAT* (pp. 145-152). Association for Computing Machinery. <https://doi.org/10.1145/3381991.3395616>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Benchmarking UAQ Solvers

Alessandro Armando
alessandro.armando@unige.it
DIBRIS, Università di Genova
Genova, Italia

Giorgia Gazzarata
giorgia.gazzarata@dibris.unige.it
DIBRIS, Università di Genova
Genova, Italia

Fatih Turkmen
f.turkmen@rug.nl
University of Groningen
Groningen, Netherlands

ABSTRACT

The User Authorization Query (UAQ) Problem is key for RBAC systems that aim to offer permission level user-system interaction, where the system automatically determines the roles to activate in order to enable the requested permissions. Finding a solution to a UAQ problem amounts to determining an optimum set of roles to activate in a given session so to obtain some permissions while satisfying a collection of authorization constraints, most notably Dynamic Mutually-Exclusive Roles (DMER) constraints. Although the UAQ Problem is NP-hard, a number of techniques to solve the UAQ problem have been put forward along with encouraging, albeit inconclusive, experimental results. We propose a methodology for designing parametric benchmarks for the UAQ problem and make a novel suite of parametric benchmarks publicly available that allows for the systematic assessment of UAQ solvers over a number of relevant dimensions. By running three prominent UAQ solvers against our benchmarks, we provide a very comprehensive analysis showing (i) the shortcomings of currently available benchmarks, (ii) the adequacy of the proposed methodology and (iii) that the reduction to PMaxSAT is currently the most effective approach to tackling the UAQ problem.

CCS CONCEPTS

• Security and privacy → Access control; Authorization.

KEYWORDS

Role-based Access Control; User Authorization Query Problem; Boolean Satisfiability

ACM Reference Format:

Alessandro Armando, Giorgia Gazzarata, and Fatih Turkmen. 2020. Benchmarking UAQ Solvers. In *Proceedings of the 25th ACM Symposium on Access Control Models and Technologies (SACMAT '20)*, June 10–12, 2020, Barcelona, Spain. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3381991.3395616>

1 INTRODUCTION

The User Authorization Query (UAQ) Problem [5] amounts to determining an optimum set of roles to activate in a given session in

order to obtain some permissions while satisfying a collection of authorization constraints, most notably Dynamic Mutually-Exclusive Roles (DMER) constraints.

The UAQ problem is key to support *permission level* user-system interaction (where the system automatically determines the roles that need to be activated in order to enable the requested permissions), as opposed to *role level* interaction (where it is the user who explicitly determines and tells the system which roles must be activated). UAQ solvers can thus help to improve the usability of most advanced access management and enforcement systems (e.g. Apache Fortress¹) offering full support to the RBAC standard [12] by relieving the user from the burden of dynamically choosing the roles to activate for the task at hand. This choice could indeed be complex as the set of roles to activate must (i) enable the execution of desired functionalities, (ii) comply with the authorization constraints prescribed by the security policy and, e.g., (iii) minimize the activation of unnecessary permissions (which represent a risk). UAQ solvers can also play an important role by offering improved support to the enforcement of RBAC policies in microservices architectures.² In this context, determining what is the most appropriate set of roles to activate for a given service is not only difficult but must be done very quickly.

In a UAQ problem the permissions requested by the user come in two sets: a lower bound P_{lb} and an upper bound P_{ub} such that $P_{lb} \subseteq P_{ub} \subseteq P$, where P is the complete set of permissions. The permissions in P_{lb} are those that *must* be granted, whereas those in $P_{ub} \setminus P_{lb}$ are additional permissions that *can* be granted. It is then possible to either minimize or maximize the number of permissions in $P_{ub} \setminus P_{lb}$ to be granted depending on which objective (safety or availability, respectively) has priority. If safety (availability) is chosen, then the number of permissions from $P_{ub} \setminus P_{lb}$ needs to be minimized (maximized, resp.). Notice that a certain degree of safety is achieved even if availability is preferred over safety, since no permission in $P \setminus P_{ub}$ can be granted. DMER constraints are of the form $DMER(\{r_1, \dots, r_m\}, t)$. They constrain activation of roles by requiring that no user can activate t or more roles in $\{r_1, r_2, \dots, r_m\}$.

Although the UAQ Problem is NP-hard [4, 5], a number of techniques to solve the UAQ problem have been put forward along with encouraging experimental results. These approaches can be broadly classified in two classes: *search-based techniques* [11, 15], whereby the problem is solved by systematically exploring a suitably defined search space, and *SAT-based techniques* [2, 10, 11, 15], i.e. techniques that leverage an encoding of the problem into either (i) a sequence of propositional satisfiability problems (SAT) [10, 11, 15] or (ii) a partial maximum propositional satisfiability problem (PMaxSAT) [2, 15] and then use SAT solvers and PMaxSAT solvers respectively as

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SACMAT '20, June 10–12, 2020, Barcelona, Spain

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7568-9/20/06...\$15.00

<https://doi.org/10.1145/3381991.3395616>

¹<https://directory.apache.org/fortress>

²<https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

workhorses to find a solution to the problem (if any). SAT-based approaches leverages the fast paced advancements achieved by lively and dedicated research communities that organize competitions³ and evaluations⁴ of state-of-the-art solvers on a yearly basis.

Most of the techniques proposed in the literature have been experimentally evaluated by running them against different benchmark problems. These benchmarks are usually parametric in some relevant dimension of the problem (e.g. number of roles, number of DMER constraints, number of requested permissions) and aim at evaluating the scalability of the proposed techniques along them. The current state of affairs is nevertheless unsatisfactory for a number of reasons. The available benchmarks do not cover (and thus do not test the solvers against) all the relevant aspects of the problem. For instance, the problems used in [11] do not consider the case where the number of roles to activate is maximized (i.e. $obj = max$) and therefore do not allow the evaluation of UAQ solvers in case of $obj = max$. Furthermore, solvers are often evaluated against problems proposed by the same authors and this does not permit to assess the relative merits of the proposed techniques.

By leveraging the asymptotic complexity analysis results given in [11], in this paper we propose a methodology for designing parametric benchmarks for the UAQ problem. As we will see, our methodology leads to benchmarks capable to (i) stress test solvers along dimensions of the problem for which no polynomial-time technique is known but also (ii) to check their effectiveness, by determining whether they efficiently solve problems that are known to be solvable in polynomial time.

By using our methodology we introduce and make publicly available a novel suite of parametric benchmarks that allows for the systematic assessment of UAQ solvers over a number of relevant dimensions. These include problems for which no polynomial-time algorithm is known as well as problems for which polynomial-time algorithms do exist. For each benchmark we indicate its purpose and the expected behavior of solvers. The suite consists of 23 parametric benchmarks grouped in 14 families: 11 benchmarks (from 7 families) have the safety objective and 12 (from 7 families) have the availability objective. For each benchmark and each considered value of the parameter the suite includes 10 different problem instances. This contribution is particularly important. In fact, the lack of adequate benchmarks makes it difficult to assess the efficiency of the different techniques used to tackle the UAQ problem. Even more so, it makes it difficult to compare their performance. By providing a complete and adequate suite of benchmarks, we are de facto offering the possibility to both assess the performance of a single solver, and compare the performance of the different solutions. [7] provided a similar contribution for the famous Symmetric and Asymmetric Traveling Salesperson Problems. This work helped to (i) reduce the publication of relatively weak TSP heuristics and (ii) motivate development of new TSP approaches.

We have used the new suite of benchmark problems as well as the benchmarks introduced in [11] to experimentally evaluate three state-of-the-art solvers: 2D-Opt-Search [11] a search-based solver, 2D-Opt-CNF [11] that combines the reduction of the UAQ Decision Problem to SAT, a state-of-the-art SAT solver and a binary search,

and AQUA [2], a SAT-based solver that implements a reduction of the UAQ Problem to PMaxSAT and uses any state-of-the-art PMaxSAT solver to tackle the problem. The experiments provide a comprehensive and comparative analysis of unprecedented breadth from which it is possible to draw a number observations:

- 2D-Opt-CNF and AQUA quickly solve all benchmark problems taken from [11]. This result indicates that this suite of benchmarks does not represent adequately the complexity of the problem and it is therefore of limited use to assess the effectiveness of the solvers.
- 2D-Opt-CNF and 2D-Opt-Search scale poorly for (most) benchmarks in our suite for which polynomial-time algorithms exist. This results is in stark contrast with the experimental results given in [11] from which one could be led to believe that these two solvers scale well in practice.
- For most benchmarks (22 out of 23), the behavior of AQUA meets the expectations that can be drawn from the known asymptotic complexity analysis results; this results is a strong indication of the validity of the methodology we propose in this paper, but it also indicates that there is still room for further investigation.
- AQUA outperforms both 2D-Opt-Search and 2D-Opt-CNF in the vast majority of the benchmark considered. This is hardly surprising since PMaxSAT solver implements sophisticated search algorithms specifically tailored to tackle optimization problems; in contrast, 2D-Opt-CNF tackles the optimization problem through a fairly naive binary search strategy and 2D-Opt-Search simply enumerates the solutions in order to find an optimum. Yet, prior to our experimental analysis, there was little empirical evidence (if any) to support this conclusion.
- We single out a benchmark, namely Pub, which can be solved efficiently but that all solvers we considered in our experiments find it difficult to tackle. We believe that further analysis on this family of problems could lead to improvements of existing solvers.

For reasons of space, in this paper we present the specifications and the experimental results of a subset of the benchmarks we generated. The interested reader can find more details [6?].

Structure of the paper. In the next section we introduce the UAQ problem. In Section 3 we provide an overview of the techniques for solving the problem. In Section 4 we present our methodology for the generation of parametric benchmark problems and introduce the new suite of benchmarks. In Section 5 we present and discuss the experimental results. In Section 6 we discuss the related work and in Section 7 we conclude the paper with some final remarks.

2 THE UAQ PROBLEM

An RBAC policy is a tuple $RP = (U, R, P, UA, PA, \geq, C)$, where U is a set of users, R a set of roles, and P a set of permissions; users are associated to roles by the *user-assignment* relation $UA \subseteq U \times R$ and roles are associated to permissions by *permission-assignment* relation $PA \subseteq R \times P$; $\geq$ is a partial order on R , modeling the hierarchy between roles, i.e. $r_1 \geq r_2$ means that r_1 is *more senior than* (has all permissions of) r_2 for $r_1, r_2 \in R$; and C is a set of *dynamic mutually exclusive role (DMER)* constraints of the form

³<http://www.satcompetition.org>

⁴<https://maxsat-evaluations.github.io/>

DMER($\{r_1, \dots, r_m\}, t$), with $t \leq m$, stating that no user can simultaneously activate t or more roles in the set $\{r_1, \dots, r_m\}$.

Let $u \in U$, we define $R_u = \{r' \in R : r \geq r' \text{ for some } r \in R \text{ such that } (u, r) \in UA\}$. A user u has permission p iff $(r, p) \in PA$ for some $r \in R_u$. We assume that RBAC policies considered are *finite*, i.e. U , R , and P have finite cardinality. Moreover, we treat permissions as if they are opaque (i.e. we do not consider the internal structure of permissions) and mutually independent, i.e. the possession of one or more permissions does not imply the possession of another permission. Let $p \in P$, we define $R_p = \{r' \in R : r' \geq r \text{ for some } r \in R \text{ such that } (r, p) \in PA\}$ and $R_{p'} = \bigcup_{p \in P'} R_p$ for any $P' \subseteq P$. Similarly, if $r \in R$ then $P_r = \{p \in P : (r', p) \in PA \text{ for some } r' \in R \text{ such that } r \geq r'\}$ and $P_{R'} = \bigcup_{r \in R'} P_r$ for any $R' \subseteq R$. Let $q \subseteq R$, then we say that q satisfies DMER($\{r_1, \dots, r_m\}, n$) iff $|\{r_1, \dots, r_m\} \cap q| < n$ and that q satisfies C iff q satisfies c for all $c \in C$.

Let S be a set of sessions and $user : S \rightarrow U$ a function that associates each session $s \in S$ with the corresponding user.

A *User Authorization Query* (UAQ) is a tuple $q = (s, P_{lb}, P_{ub}, obj)$, where $s \in S$, $P_{lb} \subseteq P_{ub} \subseteq P$, and $obj \in \{\text{any}, \text{min}, \text{max}\}$.

Definition 2.1 (UAQ Problem). The UAQ Problem for $q = (s, P_{lb}, P_{ub}, obj)$ in RP is the problem of determining a set of roles $q \subseteq R_{user(s)}$ such that (i) q satisfies C , (ii) $P_{lb} \subseteq P_q \subseteq P_{ub}$, and (iii) any other $q' \subseteq R_{user(s)}$ that satisfies C and $P_{lb} \subseteq P_{q'} \subseteq P_{ub}$ is such that $P_q \subseteq P_{q'}$ if $obj = \text{min}$, or $P_{q'} \subseteq P_q$ if $obj = \text{max}$. In case $obj = \text{any}$, any q satisfying the (i) and (ii) is a valid solution.

Definition 2.2 (UAQ Decision Problem). Let $w = (s, P_{lb}, P_{ub}, k_p)$ where $k_p \in \{\leq, \geq\} \times [0, |P_{ub} \setminus P_{lb}|]$. The UAQ Decision Problem for w in RP is the problem of determining a set of roles $q \subseteq R_{user(s)}$ such that (i) q satisfies C , (ii) $P_{lb} \subseteq P_q \subseteq P_{ub}$, and (iii.a) $|P_q \setminus P_{lb}| \leq n$ if $k_p = (\leq, n)$, or (iii.b) $|P_q \setminus P_{lb}| \geq n$ if $k_p = (\geq, n)$.

For the sake of simplicity but without loss of generality, in our work we assume that no role in R is assigned permissions that are not contained in P_{ub} , and that $P_{ub} = P$. In fact, all the permissions in $P \setminus P_{ub}$ must not be granted; therefore all the roles that activate these permissions cannot be included in a solution. These roles and the permissions in $P \setminus P_{ub}$ can thus be safely removed from the policy (in polynomial time). By doing that, we simply exclude UAQ problems that are de facto equivalent to UAQ problems that meet our assumptions.

3 SOLVING THE UAQ PROBLEM

In this section, we provide a systematic overview of search-based and SAT-based techniques for solving the UAQ problem. This will also give insights into the complexity of the problem.

3.1 Search-based Techniques

A first approach to obtain a set of roles covering the desired permissions minimally and satisfying the constraints is proposed in [16] and improved in [15]. The approach amounts to

Algorithm SearchR

- (1) enumerating all possible role activations for the user,
- (2) checking (in polynomial time as shown in [11]) whether the selected roles grant the requested permissions (i.e. fall between P_{lb} and P_{ub}), and satisfy the DMER constraints, and

- (3) keeping the best (according to the security objective considered) solution encountered, if any.

The algorithm is clearly in $O(2^{|R|})$. However, if the objective is min or any, then it is sufficient to activate at most one role per permission and this leads to the following, more efficient version of the algorithm:

Algorithm SearchRPlb

- (1) enumerate all roles $r_p \in R_p$ for each $p \in P_{lb}$,
- (2) check in polynomial time whether $S = \{r_p \in R_p : p \in P_{lb}\}$ is such that $P_S \subseteq P_{ub}$ and S satisfies the DMER constraints, and
- (3) keep the best (according to the security objective considered) solution encountered, if any.

If the objective is min or any, it is in fact possible to consider the activation of individual roles granting the permissions in P_{lb} and the algorithm is in $O(|R_P|^{|P_{lb}|})$ [11]. Instead, an alternative approach to tackling the problem for the max case is put forward in [11]:

Algorithm SearchRDMER

- (1) enumerate all sets of possible role activations $R_a = R_1 \cup \dots \cup R_n \cup R_{free}$, where $R_i \subseteq rs_i$ and $|R_i| < t_i$, for all constraints $DMER(rs_i, t_i)$ in C and $i = 1, \dots, |C|$, and $R_{free} \subseteq R$ is the set roles that do not occur in the DMER constraints (and can thus be freely activated),
- (2) check in polynomial time whether $P_{lb} \subseteq P_{R_a} \subseteq P_{ub}$ and R_a satisfies the DMER constraints.
- (3) keep the maximum sized $|P_{R_a}|$ solution encountered, if any.

We now note that for each constraint $DMER(rs_i, t_i)$ in C there are $\sum_{k=1}^t \binom{|rs_i|}{k}$ subsets R_i of rs_i such that $|R_i| < t_i$. The number of sets $R_1 \cup \dots \cup R_n$ is $\left(\sum_{k=1}^t \binom{|rs_i|}{k}\right)^{|C|}$, from which it easily follows that the enumeration of the set of roles $R_a = R_1 \cup \dots \cup R_n \cup R_{free}$ grows as $O(\widehat{rs}^{|C|\widehat{t}})$, where $\widehat{rs} = \max_{DMER(rs, t) \in C} |rs|$ and $\widehat{t} = \max_{DMER(rs, t) \in C} t$. This improves the upper bound $O(|R|^{|C|\widehat{t}})$ given in [11]. As shown in [11], the role hierarchy, $\geq$, does not contribute to the computational complexity of the problem.

3.2 SAT-based Techniques

Let $RP = (U, R, P, UA, PA, \geq, C)$ be an RBAC policy with constraints and $q = (s, P_{lb}, P_{ub}, obj)$ a UAQ query for RP . Since RP is finite (i.e. the set U of users, R of roles, and P of permissions are all finite), the UAQ problem can be tackled by leveraging SAT solvers. This can be done in a variety of ways. A first approach [11] amounts to reducing the UAQ Decision Problem to SAT and solving the optimization problem through binary search that leverage the SAT solver as an oracle for the decision problem. The second approach ([2, 14, 15]) eliminates the need for the binary search by directly encoding UAQ problems into PMaxSAT.

3.2.1 Reducing the UAQ Decision Problem to SAT. Let $w = (s, P_{lb}, P_{ub}, k_p)$ where $k_p \in \{\leq, \geq\} \times [0, |P_{ub} \setminus P_{lb}|]$. The reduction of the UAQ Decision Problem for w in RP to SAT amounts to generating a set of clauses $C_{RP}^w = C_{RP} \cup C^q$, where C_{RP} and C^w are defined

below. We assume the existence of a propositional variable $\bar{r}$ for each $r \in R$ and a propositional variable $\bar{p}$ for each $p \in P$.

C_{RP} is the smallest set of propositional clauses satisfying the following conditions.

Core RBAC.

- (1) for all $r \in R$ if $(user(s), r) \notin UA$ then $\neg \bar{r} \in C_{RP}$;
- (2) for all $p \in P$ and $r \in R$ such that $(p, r') \in PA$ with $r \geq r'$, $(\neg \bar{r} \vee \bar{p}) \in C_{RP}$;
- (3) for all $p \in P$, $(\neg \bar{p} \vee \bigvee \{\bar{r} : \text{exists } r' \in R, r \geq r', (p, r') \in PA\}) \in C_{RP}$.

It is easy to see that the number of clauses above is in $O(|R||P|)$ and the number of propositional variables in $O(|R| + |P|)$.

DMER Constraints. For all $DMER(rs, n) \in C$, a CNF of the formula $\sum_{r \in rs} \bar{r} \leq n-1$ is in C_{RP} . As shown in [13], inequalities of the form $\sum_{x \in X} x \leq n$ can be succinctly encoded into CNF with $7|X|$ clauses and $2|X|$ additional propositional variables. Thus, any DMER in C can be encoded with a number of variables and clauses in $O(|R|)$.

C^w is the smallest set of propositional clauses satisfying the following conditions.

Query.

- a unit clause $\bar{p} \in C_{RP}$ for each $p \in P_{lb}$;
- a unit clause $\neg \bar{p} \in C_{RP}$ for each $p \in P \setminus P_{ub}$;
- A CNF of the formula $\sum_{p \in P_{ub} \setminus P_{lb}} \bar{p} \bowtie n$ in C_{RP} , where $\bowtie$ is $\leq$ or $\geq$.

It can be shown that any solution to C_{RP}^w corresponds to a solution of the corresponding UAQ Decision Problem and vice versa.

3.2.2 Reducing the UAQ Problem to PMaxSAT. A PMaxSAT problem is given by a pair $\langle \mathcal{H}, \mathcal{S} \rangle$, where $\mathcal{H}$ and $\mathcal{S}$ are two finite sets of clauses, called “hard” and “soft” respectively. A solution to a PMaxSAT problem $\langle \mathcal{H}, \mathcal{S} \rangle$ is any truth-value assignment to the variables in $\mathcal{H}$ and $\mathcal{S}$ that satisfies all clauses in $\mathcal{H}$ and the maximum number of clauses in $\mathcal{S}$.

A UAQ Problem for $q = (s, P_{lb}, P_{ub}, obj)$ can be reduced to a PMaxSAT problem $\langle \mathcal{H}_{RP}^q, \mathcal{S}^q \rangle$, where (i) $\mathcal{H}_{RP}^q$ is obtained from C_{RP} (as defined in Section 3.2.1) by adding a unit clause $\bar{p}$ for each $p \in P_{lb}$ and a unit clause $\neg \bar{p}$ for each $p \in P \setminus P_{ub}$; and (ii) $\mathcal{S}^q$ comprises a unit clause $\neg \bar{p}$ if $obj = \min$ and a unit clause $\bar{p}$ if $obj = \max$ for each $p \in P_{ub} \setminus P_{lb}$. No soft clauses are included if $obj = \text{any}$.

It can be shown that any solution to C_{RP}^w (C_{RP}^q) corresponds to a solution of the corresponding UAQ Decision Problem (UAQ Problem, resp.) and vice versa. It can be readily seen that for both reductions the number of clauses is in $O(|R||P|)$ and the number of propositional variables is in $O(|R| + |P|)$.

4 BENCHMARKS

Designing benchmarks suitable for the systematic assessment of UAQ solvers is not easy. A common approach [2, 8, 11, 15] is to focus on families of problems that are parametric on some aspects of the problem that may contribute to its complexity. All other aspects are either set to a predefined, constant value or are randomly chosen in a given interval or according to some criterion. By running a solver against the instances corresponding to increasing values of the parameter, it is thus possible to obtain an estimation of how the solver scales along the dimension represented by the parameter.

Unfortunately, the adequacy of the benchmarks proposed in the literature is seldom discussed. For instance, as we will see, some benchmarks have been reported to be solved efficiently by the proposed solvers, e.g., in linear time along some parameter for which no polynomial-time algorithm is known. When this is the case, the likely explanation is that the benchmarks do not represent the complexity of the problem. (The alternative being that the solver used in the experiments improves over the known complexity results.)

The sheer number of elements that contribute to the definition of the UAQ problem complicates the selection of the parameters. The elements characterizing the RBAC policy include the number of roles $|R|$, the number of permissions $|P|$, the number of DMER constraints $|C|$ as well as their specific features (e.g. $\widehat{rs}$ and $\widehat{t}$). One may even consider features of the PA relation, such as the maximum number of roles that contain any given permission (referred as $\widehat{R}_p = \max_{p \in P} |R_p|$ in Section 3). The components of the query also contribute to the complexity of the problem. These include the security objective (any, min, max), the number of requested permissions that *must* be granted, i.e. $|P_{lb}|$, and the number of requested permissions that *can* be granted, i.e. $|P_{ub}|$.

In previous work (see, e.g., [2, 8, 11, 15]), various benchmark problems parametric in any of these aspects have been put forward. To the best of our knowledge, the most extensive collection of parametric benchmarks so far is presented in [11]⁵. The suite consists of the benchmarks *roles*, *d*, *rolesPerConstr*, *t*, and *plb*, which are parametric in $|R|$, $|C|$, $\widehat{rs}$, $\widehat{t}$, and $|P_{lb}|$ respectively. While the benchmarks in [11] are a first attempt to provide a comprehensive evaluation along a number of significant dimensions, they still suffer from the following shortcomings:

- (1) it is not always clear if and how these benchmarks represent the complexity of the UAQ problem;
- (2) only the optimization objective min is considered, and they are therefore not suitable for evaluating the performance of the solvers when different optimization objectives, most notably max, are considered.

In order to overcome these limitations, we introduce a methodology that can be used both to evaluate existing benchmarks and to guide the design of new ones. The complexity results introduced in Section 3 can be used to both validate and guide the design of benchmarks by leveraging the following observations:

- If a family of UAQ problems is known to be solvable in polynomial time, then by running a solver against this family of problems we can check how the performance of the solver compares with that of the known algorithms: if the time spent by the solver grows, e.g., exponentially as the size of the problem increases, then the solver is clearly inefficient against this family of problems.
- Dually, if a family of UAQ problems is known to be NP-hard, but the time spent by any solver has a polynomial growth as the size of the problem increases, then, assuming that $P \neq NP$, this means that the family of the problems (i.e. the benchmark) considered does not represent adequately the complexity of the problem.

⁵Available at <https://ece.uwaterloo.ca/~tripunit/uaq/Tested-Code.tgz>.

Table 1: Benchmark specifications for $obj = \min$

Name	$ R $	$ P_{ub} $	$\widehat{R}_P$	$ P_{lb} $
<i>Plb_bigR</i>	200	400	5	5..50
<i>Plb_smallR</i>	10	400	5	5..50
<i>R_bigPlb</i>	10..100	400	5	100
<i>R_smallPlb</i>	10..100	400	5	2
<i>RPhat_bigPlb</i>	200	400	2..12	10
<i>RPhat_medPlb</i>	200	400	2..12	4
<i>RPhat_smallPlb</i>	200	400	2..12	1

The algorithms introduced in Section 3.1 play a key role in the proposed methodology. In fact, they provide upper bounds on the asymptotic growth rate for UAQ solvers which, to the best of our knowledge, are the best upper bounds currently available in the literature. Yet, it must be noted that the proposed methodology will yield different, improved results as soon as new complexity results will become available. Thus the benchmarks proposed and used in this paper could be improved consequently.

The methodology described above can be used to understand and validate the benchmarks presented in [11]. Comparing the complexity of Algorithm SearchR ($O(2^{|R|})$) and Algorithm SearchRPlb ($O(\widehat{R}_P^{|P_{lb}|})$), it appears that Algorithm SearchRPlb (asymptotically) outperforms Algorithm SearchR over all benchmarks. We then observe that the value of $|P_{lb}|$ is fixed in all benchmarks except *plb*. Algorithm SearchRPlb is therefore insensitive to the value of the respective parameter (R for *roles*, $|C|$ for *d*, etc.) and so should be any reasonably efficient solver when applied to these problems. All benchmark problems but *plb* can thus be used to check whether UAQ solvers are as effective as Algorithm SearchRPlb as the value of the respective parameter increases. Benchmark *plb* is instead parametric in $|P_{lb}|$ and we therefore expect the solving time of Algorithm SearchRPlb (and of any other solver) to increase exponentially as $|P_{lb}|$ increases.

Driven by the our methodology, we propose two new families of parametric UAQ problems, one with $obj = \min$ and one with $obj = \max$. The benchmarks with $obj = \min$ are summarized in Table 1:

- *Plb_bigR* and *Plb_smallR* are both parametric in $|P_{lb}|$. *Plb_bigR* is designed to stress test solvers for increasing values of $|P_{lb}|$: for large values of $|R|$ (here set to 200), the best known algorithm (Algorithm SearchRPlb) is exponential in $|P_{lb}|$ and thus we expect any solver to exhibit the same behavior. *Plb_smallR* is instead designed to check the effectiveness of solvers: Algorithm SearchR is in $O(2^{|R|})$ and thus we know that the problem can be solved efficiently for sufficiently small values of $|R|$ (here set to 10).
- *R_bigPlb* and *R_smallPlb* are both parametric in $|R|$ and are dual to *Plb_bigR* and *Plb_smallR* respectively. *R_bigPlb* is designed to be used to stress test solvers for increasing values of $|R|$: for large values of $|P_{lb}|$ (here set to 100), the best known algorithm (i.e. Algorithm SearchR) is exponential in $|R|$ and thus we expect any solver to exhibit the same behavior. *R_smallPlb* can be used to check the efficiency of solvers:

Table 2: Benchmark specifications for $obj = \max$

Name	$ R $	$ P_{ub} $	$\widehat{R}_P$	$ C $	$\widehat{r}_s$	$\widehat{t}$	$ P_{lb} $
<i>R_bigCt</i>	10..100	400	5	50	8	3	10
<i>R_smallCt</i>	10..100	400	5	5	3	2	10
<i>C_bigR</i>	200	400	5	10..100	8	3	10
<i>C_smallR</i>	10	400	5	10..100	8	3	10
<i>that_bigR</i>	1000	1000	1	50	20	2..12	10
<i>that_smallR</i>	20	400	5	10	12	2..12	10
<i>rshat_bigCt</i>	200	400	5	10	5..50	3	10
<i>rshat_medCt</i>	200	400	5	3	5..50	3	10
<i>rshat_smallCt</i>	200	400	5	1	5..50	3	10

Algorithm SearchRPlb is in $O(\widehat{R}_P^{|P_{lb}|})$ and thus we know that the problem can be solved efficiently for sufficiently small values of $|P_{lb}|$ (here set to 2).

- *RPhat_bigPlb*, *RPhat_medPlb*, *RPhat_smallPlb* are parametric in $\widehat{R}_P$ and can be used to check the effectiveness of solvers: for big values of $|R|$ (here set to 200) the best algorithm is Algorithm SearchRPlb which is in $O(\widehat{R}_P^{|P_{lb}|})$. We thus know that the problem can be solved efficiently for sufficiently small values of $|P_{lb}|$ (here set to 10, 4 and 1 respectively). Note that $|P_{lb}|$ is the degree of the polynomial and therefore the time spent by the solver may differ significantly (for the values of $|P_{lb}|$ considered) as $\widehat{R}_P$ increases.

All the instances presented in the benchmarks above have no DMER constraints. The latter, in fact, do not contribute to the UAQ problem complexity when $obj = \min$ [10].

The benchmarks with $obj = \max$ are summarized in Table 2:

- *R_bigCt* and *R_smallCt* are parametric in $|R|$. *R_bigCt* is designed to stress test solvers for increasing values of $|R|$: when $|C|$ and $\widehat{t}$ are such that $|C|\widehat{t}$ is sufficiently large (here set to 150), the best algorithm (Algorithm SearchR) is exponential in $|R|$ and thus we expect any solver to exhibit the same behavior. *R_smallCt* is instead designed to check the effectiveness of solvers: when $|C|$ and $\widehat{t}$ are such that $|C|\widehat{t}$ is sufficiently small (here set to 10), Algorithm SearchRDMER can efficiently solve the problems independently from the size of $|R|$ and any efficient solver should do the same.
- *C_bigR* and *C_smallR* are parametric in $|C|$. *C_bigR* can be used to stress-test solvers for increasing values of $|C|$: for large values of $|R|$ (here set to 200), Algorithm SearchRDMER is to be preferred to Algorithm SearchR. Since Algorithm SearchRDMER is exponential in $|C|$, we expect solvers to exhibit the same behavior. *C_smallR* can be used to check the efficiency of solvers: Algorithm SearchR is in $O(2^{|R|})$ and thus the problem can be solved efficiently for sufficiently small values of $|R|$ (here set to 10).
- *that_bigR* and *that_smallR*, parametric in $\widehat{t}$, are analogous to the previous case. We then expect that solvers exhibit exponential growth over *that_bigR*, while *that_smallR* can be solved efficiently.
- *rshat_smallCt*, *rshat_medCt* and *rshat_bigCt* are parametric in $\widehat{r}_s$. For big values of $|R|$ (here set to 200), Algorithm4 outperforms Algorithm SearchR. Algorithm SearchRDMER

is in $O(\widehat{rs}^{|\widehat{C}|\widehat{t}})$ and these problems can thus be solved efficiently for sufficiently small values of $|\widehat{C}|\widehat{t}$ (here set to 3). *rshat_smallCt* can then be used to check the efficiency of solvers. *rshat_medCt* and *rshat_bigCt* can be used to see how the values of $|\widehat{C}|\widehat{t}$ (here set to 9 and 30) affect the complexity of the problem.

Notice that we do not include benchmarks parametric in the “size” of the role hierarchy since, as already pointed out in Section 3, it does not contribute to the complexity of the UAQ problem. The benchmarks in Table 1 and in Table 2 are obtained by randomly generating a RBAC policy with the specified number of roles and permissions (ensuring that each permission is assigned to exactly $\widehat{R_P}$ roles), the specified number of DMER constraints of the form $\text{DMER}(\{r_1, \dots, r_m\}, t)$ with $m = \widehat{rs}$, $t = \widehat{t}$ and a query with P_{lb} with $|P_{lb}|$ randomly selected permissions and $P_{ub} = P$.

5 EXPERIMENTAL RESULTS

In Section 4, we presented an overview of the existing benchmarks and introduced a methodology to evaluate and even guide the design of benchmarks parametric with the relevant aspects of the UAQ problem. In this section, we first discuss the benchmarks presented in [11] and then experimentally compare the performance of the following solvers against those benchmarks and the ones presented in Table 1:

- 2D-Opt-Search [11]: a search-based solver leveraging the FPP result;
- 2D-Opt-CNF [11]: a SAT-based solver that leverages the reduction of the UAQ Decision Problem to SAT, zChaff [1] as SAT solver a state-of-the-art SAT solver and a two-dimensional binary search to solve UAQ problems;
- AQUA [2]: a SAT-based solver that employs a reduction of the UAQ Problem to PMaxSAT and employs any state-of-the-art PMaxSAT solver to tackle the problem. In the experiments presented in this paper we used the Loandra PMaxSAT solver [3].

All the solvers mentioned above support joint optimization of permissions and roles. However, since we focus on problems where the optimization objective is limited to permissions only, we disabled the role optimization in 2D-Opt-Search and 2D-Opt-CNF.

In the second set of experiments, we focus our attention to *obj = max* case and evaluate the performance of AQUA against the benchmarks in Table 2.

For each benchmark problem we generated 10 instances and ran the solvers against them. All data points in our plots represent the median value of the time spent by the corresponding solver. The experiments have been conducted on a PC with 2 64-bit Intel Xeon CPU X7350 (8 core) @ 2.93GHz and 47 GB RAM running Linux (Ubuntu 16.04.5 LTS).

For all the benchmarks we set the timeout to 600 seconds, after which the solver execution is interrupted and the problem is labeled as “skipped”. We omitted the timed-out solvers from the plots.

5.1 Benchmarks from [11]

In the first experiment, we evaluate the benchmarks provided by [11] by running AQUA over them. Comparing the complexities of the algorithms presented in Section 3.1, it appears that Algorithm

SearchRPlb should outperform the other algorithms over all the benchmarks. Since Algorithm SearchRPlb is in $O(\widehat{R_P}^{|P_{lb}|})$ and the value of $|P_{lb}|$ is fixed in *roles*, *d*, *rolesPerConstr*, and *t*, we expect that AQUA is insensitive to the variation of the respective parameters. This expectation is confirmed by the experimental results: in particular, AQUA solved all the instances contained in *roles*, *d*, *rolesPerConstr*, and *t* in less than 20 milliseconds. Instead, since *plb* is parametric in $|P_{lb}|$, AQUA should exhibit an exponential growth over it. Contrary to our expectation, AQUA could solve all the problems contained in the benchmark in a fraction of a second (at most 4 milliseconds) even for very large values of $|P_{lb}|$. This result means that the benchmark *plb* does not represent the complexity of the UAQ problem on varying numbers of $|P_{lb}|$. From the experiment, it follows that the benchmarks provided by [11] cannot be used to stress test solvers. They can nevertheless be used to evaluate their efficiency.

We then ran 2D-Opt-Search and 2D-Opt-Search over the benchmarks. AQUA outperformed them over all the benchmarks. In particular, 2D-Opt-CNF exhibited a behavior similar to AQUA, however there is always a distance of two orders of magnitude between the median solving time of the two solvers. Also 2D-Opt-Search over *roles*, *d*, *t*, and *rolesPerConstr* behaved like AQUA, always outperforming 2D-Opt-CNF. Instead, the time spent by 2D-Opt-Search over *plb* is exponential.

5.2 Our benchmarks with *obj = min*

We now turn our attention to our benchmark proposals where *obj = min* (cf. Table 1). Figure 1 depicts the results of their experimental evaluation:

- The time spent by AQUA to solve *Plb_bigR* and *Plb_smallR* benchmarks meets the expectations. For *Plb_bigR*, the plot has a clear exponential growth, whereas for *Plb_smallR*, AQUA grows very slowly and even (slowly) decreases for $|P_{lb}| > 20$.
- Similar considerations hold for *R_bigPlb* and *R_smallPlb*.
- The results for *RPhat_bigPlb*, *RPhat_medPlb*, *RPhat_smallPlb* are of particular interest. As pointed out in Section 3, Algorithm SearchRPlb is in $O(\widehat{R_P}^{|P_{lb}|})$ and if $|P_{lb}|$ is bounded, then it can solve the problem in time which grows as a polynomial of degree $|P_{lb}|$. It can be noted that AQUA quickly solves all problems in *RPhat_smallPlb* (for which $|P_{lb}| = 1$), while the plot is considerably more steep for *RPhat_bigPlb* (for which $|P_{lb}| = 12$). This is consistent with the growth of polynomials of degree 1 and 12 respectively. The plot for *RPhat_smallPlb* (for which $|P_{lb}| = 4$) represents an intermediate situation.

The presented experimental results confirm that the benchmarks of Table 1 for problems with *obj = min* is complete and satisfactory. It consists of hard benchmarks (*Plb_bigR* and *R_bigPlb*) which represent the complexity of the UAQ problem and can be used to stress test the solvers, and easy benchmarks (*Plb_smallR*, *R_smallPlb*, *RPhat_bigPlb*, *RPhat_medPlb*, and *RPhat_smallPlb*), which can be used to evaluate the solvers’ performance. Figure 1 then also shows the results of the solvers’ performance evaluation:

Figure 1: Performance of 2D-Opt-Search, 2D-Opt-CNF and AQUA on the benchmarks of Table 1

- The behavior of 2D-Opt-CNF on *Plb_smallR* is similar to that of AQUA, but when applied to *Plb_bigR* it reaches the timeout even for the smallest value of $|P_{lb}|$ (i.e. 5). 2D-Opt-Search can solve *Plb_smallR* and *Plb_bigR* only for the smallest value of $|P_{lb}|$.
- Similar considerations hold for *R_bigR* and *R_smallR*. Notice that 2D-Opt-Search performs remarkably well for *R_smallR*, but it reaches the timeout even for the smallest instance of *R_bigR* (i.e. 10).
- The behavior of 2D-Opt-Search on *RPhat_smallPlb* and *RPhat_medPlb* is similar to that of AQUA, but it reaches the timeout for *Plb_bigR* for a small value of $\widehat{R}_p$ (i.e. 2). Instead, 2D-Opt-CNF reaches the timeout for all the instances contained in the three benchmarks.

We can then conclude that the approach used by AQUA seems to allow an important reduction of the solving time compared to 2D-Opt-Search and 2D-Opt-CNF over almost all the benchmarks of Table 1.

5.3 Our benchmarks with $obj = max$

Figure 2 presents the experimental evaluation of the benchmarks with the optimization objective set to max (cf. Table 2) obtained by running AQUA against them. From the experimental results we can draw the following observations:

- The time spent by AQUA to solve *R_bigCt* and *R_smallCt* meets the expectations. For *R_bigCt* the plot has a clear exponential growth behaving like Algorithm SearchR, whereas for *R_smallCt* the plot is almost constant, behaving like Algorithm SearchRDMER, whose complexity does not depend on $|R|$.
- The time spent by AQUA to solve *C_bigR* and *C_smallR* meets the expectations. For *C_bigR* the plot has a clear exponential growth, whereas for *C_smallR* the solving time is around 100 ms independently from $|C|$.
- As expected, the median solving time of AQUA over *that_bigR* grows exponentially like Algorithm SearchRDMER, while AQUA quickly solves the instances contained in *that_smallR*, behaving like Algorithm SearchR.

- The plot for *rshat_smallCt* show that the problems can be solved quickly as $\widehat{r}_s$ increases as long as the value of $|C|\widehat{t}$ is sufficiently small (3 in this case). The plots for *rshat_medCt* and *rshat_bigCt* show the effect of larger values of $|C|\widehat{t}$ which is consistent with the growth of polynomials of degree 9 and 30 respectively.

From the evaluation, we then conclude that the benchmarks of Table 2 is complete and satisfactory for $obj = max$. In particular, *R_bigCt*, *C_bigR*, *that_bigR*, and *rshat_bigCt* are hard benchmarks representing the complexity of the UAQ problem and then can be used to stress test the solvers; on the contrary, *R_smallCt*, *C_smallR*, *that_smallR*, *rshat_medCt*, and *rshat_smallCt* are easy benchmarks and can be used to evaluate the solvers' performance.

We do not show the results of the evaluation of the efficiency of 2D-Opt-Search and 2D-Opt-CNF: 2D-Opt-Search may lead to sub-optimal results, thus the comparison with AQUA would not be fair⁶; instead, it is not clear if 2D-Opt-CNF can tackle problems with $obj = max$. In fact, we have evidence of many inconclusive solutions (neither satisfiable nor unsatisfiable).

6 RELATED WORK

As UAQ is a central problem in RBAC systems, significant effort [8–10, 14] has been put to develop techniques for tackling it efficiently and to understand its underlying complexity. To our knowledge, [5] is the first paper that discusses UAQ where authors show that the complexity of finding minimal set of roles to be activated in a session that covers the permissions requested by the user is NP-complete. While they analyze UAQ in the presence of complex role hierarchies, they do not consider the constraint types (e.g. mutual exclusion of roles) available in RBAC. The more generic form for the UAQ problem where there are lower (P_{lb}) and upper (P_{ub}) bound permissions has been first proposed in [15].

The formal complexity analysis of different UAQ problem classes is also a point of strong scientific interest. For instance, [4] shows the complexities of UAQ problem by reducing it to a special case of set covering problem namely, container optimization. The reduction shows that the UAQ problem is NP-hard when $obj = min$, while it is in P when $obj = max$. However, the paper does not consider UAQ

⁶For more details, see [?].

Figure 2: Performance of AQUA on the benchmarks of Table 2

problems with dynamic constraints as we do in this paper. Moreover, these works did not either conduct any experimental analysis or employed extremely simple RBAC instances to evaluate their proposals. Then, [11] shows that the complexity result presented in [4] for the case $obj = \max$ is different when there are constraints in the UAQ instance. More specifically, the authors show that there is an upper bound (NP) for the general UAQ problem and the case $obj = \max$ is intractable if the UAQ instances have constraints.

7 CONCLUSIONS

In this paper, we presented a systematic overview of the computational complexity, existing algorithms and available benchmarks pertinent to the UAQ problem. Our analysis of the currently available benchmarks revealed that they are inadequate to analyze the inherent complexity of different UAQ problem classes. We then proposed a methodology to generate UAQ benchmarks starting from the known complexity results and used them to evaluate the state-of-the-art UAQ solvers. Our experimental results do not only show the effectiveness of the solvers over various UAQ problem classes but also the impact of the chosen parameter in the overall performance. As future work we would like to investigate the generation for benchmarks of joint optimization.

REFERENCES

- [1] [n.d.]. *zChaff*. <http://www.princeton.edu/~chaff/zchaff.html>
- [2] F. Turkmen A. Armando, G. Gazzarata. [n.d.]. Benchmarking UAQ Solvers. https://github.com/GioGazza/uaq_prolem/tree/master/publications. Extension of paper accepted at SACMAT 2020.
- [3] Jeremias Berg, Emir Demirovic, and Peter J. Stuckey. 2019. Core-Boosted Linear Search for Incomplete MaxSAT. In *Integration of Constraint Programming, Artificial Intelligence, and Operations Research - 16th International Conference, CPAIOR 2019, Thessaloniki, Greece, June 4-7, 2019, Proceedings*. 39–56.
- [4] Liang Chen and Jason Crampton. 2009. Set covering problems in role-based access control. In *Proceedings of the 14th European conference on Research in computer security (ESORICS'09)*. 689–704.
- [5] S. Du and J. B. D. Joshi. 2006. Supporting authorization query and inter-domain role mapping in presence of hybrid role hierarchy. In *SACMAT*. 228–236.
- [6] Giorgia Gazzarata. 2020. *Extensions and Experimental Evaluation of SAT-based solvers for the UAQ problem*. Ph.D. Dissertation. University of Genova (Italy). https://github.com/GioGazza/uaq_prolem/tree/master/publications
- [7] John Larusic and Abraham P. Punnen. 2014. The asymmetric bottleneck traveling salesman problem: Algorithms, complexity and empirical analysis. *Comput. Oper. Res.* 43 (2014), 20–35.
- [8] Jianfeng Lu, James B. D. Joshi, Lei Jin, and Yiding Liu. 2015. Towards complexity analysis of User Authorization Query problem in RBAC. *Computers & Security* 48 (2015), 116–130.
- [9] Jianfeng Lu, Yun Xin, Zhao Zhang, Hao Peng, and Jianmin Han. 2018. Supporting user authorization queries in RBAC systems by role-permission reassignment. *Future Generation Comp. Syst.* 88 (2018), 707–717.
- [10] Nima Mousavi. 2014. *Algorithmic Problems in Access Control*. Ph.D. Dissertation. University of Waterloo, Canada.
- [11] Nima Mousavi and Mahesh V. Tripunitara. 2012. Mitigating the Intractability of the User Authorization Query Problem in Role-Based Access Control (RBAC). In *NSS*. 516–529.
- [12] National Institute of Standards and Technology (NIST). 2004. Role-Based Access Control. *American National Standards Institute, Inc.* (2004).
- [13] C. Sinz. 2005. Towards an Optimal CNF Encoding of Boolean Cardinality Constraints. In *Principles and Practice of Constraint Programming (CP)*. 827–831.
- [14] Fatih Turkmen. 2012. *Exploring Dynamic Constraint Enforcement and Efficiency in Access Control*. Ph.D. Dissertation. University of Trento, Italy.
- [15] G. T. Wickramaarachchi, W. H. Qardaji, and N. Li. 2009. An efficient framework for user authorization queries in RBAC systems. In *SACMAT*. 23–32.
- [16] Y. Zhang and J. B. D. Joshi. 2008. UAQ: a framework for user authorization query processing in RBAC extended with hybrid hierarchy and constraints. In *SACMAT*. 83–92.